

Thinking about Acting:

Logical Foundations for Rational Decision Making

John L. Pollock
Department of Philosophy
University of Arizona
Tucson, Arizona 85721
pollock@arizona.edu
<http://www.u.arizona.edu/~pollock>

Abstract

The objective of this book is to produce a theory of rational decision making for realistically resource-bounded agents. My interest is not in “What should I do if I were an ideal agent?”, but rather, “What should I do given that I am who I am, with all my actual cognitive limitations?”

The book has three parts. Part One addresses the question of where the values come from that agents use in rational decision making. The most common view among philosophers is that they are based on preferences, but I argue that this is computationally impossible. I propose an alternative theory somewhat reminiscent of Bentham, and explore how human beings actually arrive at values and how they use them in decision making.

Part Two investigates the knowledge of probability that is required for decision-theoretic reasoning. I argue that subjective probability makes no sense as applied to realistic agents. I sketch a theory of objective probability to put in its place. Then I use that to define a variety of causal probability and argue that this is the kind of probability presupposed by rational decision making. So what is to be defended is a variety of causal decision theory.

Part Three explores how these values and probabilities are to be used in decision making. In chapter eight, it is argued first that actions cannot be evaluated in terms of their expected values as ordinarily defined, because that does not take account of the fact that a cognizer may be unable to perform an action, and may even be unable to try to perform it. An alternative notion of “expected utility” is defined to be used in place of expected values. In chapter nine it is argued that individual actions cannot be the proper objects of decision-theoretic evaluation. We must instead choose plans, and select actions indirectly on the grounds that they are prescribed by the plans we adopt. However, our objective cannot be to find plans with maximal expected utilities. Plans cannot be meaningfully compared in that way. An alternative, called “locally global planning”, is proposed. According to locally global planning, individual plans are to be assessed in terms of their contribution to the cognizer’s “master plan”. Again, the objective cannot be to find master plans with maximal expected utilities, because there may be none, and even if they are, finding them is not a computationally feasible task for real agents. Instead, the objective must be to find good master plans, and improve them as better ones come along. It is argued that there are computationally feasible ways of doing this, based on defeasible reasoning about values and probabilities.

TABLE of CONTENTS

Chapter One: Rational Choice and Classical Decision Theory

1. Rational Cognition
2. Ideal Rationality and Real Rationality
3. Human Rationality and Generic Rationality
4. Decision Making
5. Classical Decision Theory and the Optimality Prescription

Part One: Values

Chapter Two: Evaluative Cognition and The Evaluative Database

1. The Doxastic/Conative Loop
2. Preference Rankings
3. Analog Representations of Values
 - 3.1 Analog Representations of Quantities
 - 3.2 Q&I Modules
 - 3.3 Numbers
 - 3.4 Decision-Theoretic Reasoning
 - 3.5 Theorizing about Practical Cognition
4. Conclusions

Chapter Three: Evaluative Induction

1. The Need for Evaluative Induction
2. Human Conative States
 - 2.1 Feature Likings and Desires
 - 2.2 State Likings
3. Evaluative Induction
 - 3.1 Cumulative State Liking
 - 3.2 Correcting State Liking
 - 3.3 Appropriate Measures of State Liking
 - 3.4 The Evaluative Database
 - 3.5 Conative Conditioning
 - 3.6 Evaluative Induction
4. Rational Decision Making
5. Liking Activities
6. Evaluating the Human Cognitive Architecture
7. State Liking
 - 7.1 Is State Liking the only Value?
 - 7.2 Conative Dispositions
 - 7.3 The Effects of Feature Likings on State Likings
 - 7.4 Fine-Tuning State Liking
 - 7.5 "If it feels good, keep doing it"
 - 7.6 A Puzzle
8. Conclusions

Chapter Four: Feature-Based Evaluative Cognition

1. The Database Calculation
2. Justifying the Database Calculation
3. Feature-Based Evaluative Cognition
 - 3.1 Defining Expected Values
 - 3.2 Computing Marginal Expected Values
4. Building a Cognitive Agent

Part Two: Probabilities

Chapter Five: Subjective Probabilities

1. Two Kinds of Probabilities
2. Subjective Probabilities and Degrees of Belief
3. Belief Simpliciter
4. Subjective Expected Utility Theory
5. Rational Decision Making
6. Do Subjective Probabilities Exist
7. Deriving the Optimality Prescription from Rationality Constraints
8. Subjective Probabilities from Epistemology
 - 8.1 Do Degrees of Justification Conform to the Probability Calculus?
 - 8.2 Do Degrees of Warrant Conform to the Probability Calculus?
 - 8.3 Probabilism and Reasoning
 - 8.4 Conjunctivitis
 - 8.5 Conclusions
9. A Return to Objective Probabilities

Chapter Six: Objective Probabilities

1. Physical Probabilities and Relative Frequencies
2. Empirical Theories
3. Nomic Probability
 - 3.1 The Calculus of Nomic Probabilities
 - 3.2 The Statistical Syllogism
 - 3.3 Statistical Induction
 - 3.4 Non-Classical Direct Inference
4. Mixed Physical/Epistemic Probabilities
 - 4.1 Direct Inference
 - 4.2 Reducing Classical Direct Inference to Non-Classical Direct Inference
 - 4.3 Betting the Farm
5. Conclusions

Chapter Seven: Causal Probabilities

1. Causal Decision Theory
2. Probabilistic Causation

3. Skyrms and Lewis
4. Defining Causal Probability
5. Conditional Policies and Conditional Causal Probability
6. **C-PROB_A** and **K-PROB_A**
 - 6.1 Causation and Time
 - 6.2 The Markov Condition
 - 6.3 Causal Decision Theory and Evidential Decision Theory
7. Computing Causal Probabilities
8. Computing Conditional Causal Probabilities
9. Simplifying the Computation Defeasibly
10. Conclusions

Part Three: Decisions

Chapter Eight: Rational Choice and Action Omnipotence

1. Actions and the Optimality Prescription
2. Action Omnipotence
3. Restricting the Scope of the Optimality Prescription
 - 3.1 High-Level Actions
 - 3.2 Acts and Actions
 - 3.3 Basic Actions and Action Omnipotence
 - 3.4 Deciding
 - 3.5 Willing
 - 3.6 Trying
 - 3.7 Giving Up on Infallibly Performable Actions
4. Expected Utility
5. Conditional Policies and Expected Utilities
6. Computing Expected-Utilities
7. Higher-Order Decisions
8. Conclusions

Chapter Nine: Plans and Decisions

1. Against Optimality
2. The Logical Structure of Practical Deliberation
 - 2.1 Deciding-Whether and Deciding-Which
 - 2.2 Strong and Weak Competition
 - 2.3 Boolean Combinations
 - 2.4 Universal Plans
 - 2.5 The Search for Alternatives — Resurrecting (OP)
3. Groups of Actions
4. Actions and Plans
5. Choosing Between Plans
6. AI Planning Theory: The Real World vs. Toy Problems
7. When is a Plan a Good One?
8. Locally Global Planning
9. Conclusions

Chapter Ten: Plans and Their Expected Utilities

1. Linear Plans
2. Linear Policies
3. Nonlinear Plans
4. Conditional Plans
5. An Example
6. Conclusions

Chapter Eleven: Locally Global Planning

1. The Theory
2. Incremental Decision-Theoretic Planning
3. Goal-Directed Planning
4. Presumptively Additive Expected Utilities
5. Finding and Repairing Decision-Theoretic Interference
6. Conclusions

Appendix: The Theory of Nomic Probability

1. Introduction
2. Computational Principles
3. The Statistical Syllogism
4. Direct Inference and Definite Probabilities
5. Indefinite Probabilities and Probability Distributions
6. Induction
7. Conclusions

Bibliography